EXHIBIT D

DECLARATION OF DR. BERNARD SISKIN

I, Dr. Bernard Siskin, hereby state as follows:

1. I am over the age of eighteen and am competent to make this Declaration. I have personal knowledge of the matters set forth herein.

2. I am the founder and director of BLDS, LLC.

3. BLDS was founded in 2011. It provides statistical and economic analysis to clients such as law firms, companies, government entities, and others.

4. Among the core services BLDS provides is statistical analysis of policies' and practices' discriminatory effects and identification of available alternatives that would ameliorate any discriminatory effects. BLDS is among the national leaders in that field. Its work, and the work of BLDS's various experts before they joined BLDS, has been cited repeatedly in judicial opinions and has formed the basis of many regulatory actions and corporate decisions and policies.

5. I have worked in this field since long before I founded BLDS. I have been involved in many projects that helped identify, reduce, or eliminate policies' unnecessary discriminatory effects, including in the making of home loans and other policies affecting the availability of housing and credit. I and BLDS have been retained to analyze the discriminatory effects of policies by numerous governmental and private organizations including, but not limited to, financial institutions, including large banks, the Consumer Financial Protection Bureau, the Federal Trade Commission, the Third Circuit Task Force on Race and Gender Equality in the Courts, the Equal Employment Opportunity Commission, the Civil Rights

Division of the United States Justice Department, the Office of Federal Contract Compliance, the Federal Bureau of Investigation, the Federal National Mortgage Association, the Federal Home Loan Mortgage Corporation, and various states and municipalities.

6. The disparate impact analyses that I and others at BLDS have performed have made many housing and lending policies less discriminatory and more inclusive. For example, I worked with Fannie Mae and Freddie Mac (the "GSEs") on fair lending analyses of their automated underwriting systems. These systems, which use algorithms to evaluate the riskiness of home loans, made lending decisions quicker and eliminated much of the discretion that had contributed to racially discriminatory lending decisions. But these models also risked causing unnecessary disparities adverse to protected classes. With our assistance, the GSEs were able to refine their automated underwriting systems, finding alternative policies that are comparably effective at evaluating creditworthiness while reducing disparate impacts that would have disfavored borrowers on the basis of protected classes unnecessarily.

7. BLDS has performed and continues to perform similar services for private institutions, including large financial institutions that make mortgage loans. These institutions hire BLDS to perform fair lending analyses in part because they are concerned about liability risks under the traditional disparate-impact standard. With that standard in place, they face risk if their policies or practices cause discriminatory effects and they fail to adopt alternatives that have a less discriminatory result while still accomplishing legitimate ends. Our analyses help them avoid that risk.

8. Without the risk of traditional disparate-impact liability—including the threat of liability for failing to adopt less discriminatory alternatives—all these institutions would have very different incentives. They would not have the same incentive to assess whether their

policies and practices cause disproportionate adverse effects on protected classes or whether less discriminatory alternatives existed, nor to retain me and other BLDS analysts to help them do so. All these industry actions described above took place against the backdrop of interest by regulators and private litigators in challenging discriminatory lending practices. Without that regulatory backdrop, lenders would be far less interested in our disparate-impact analyses.

9. Models and algorithms have been used for housing-related decisions for years, with the most obvious examples being automated credit underwriting, line assignment and pricing. In the last two decades, many lenders and others whose actions affect the availability of housing have increasingly adopted robust compliance functions that monitor policies for disparate impact and seek alternatives with less discriminatory results. Many regularly retain the assistance of BLDS to perform and assist with these functions.

10. To meet the needs of this market created by the traditional disparate-impact doctrine, BLDS has created sophisticated methodologies and proprietary algorithms. These tools allow BLDS to assess whether policies have a disparate impact and then identify whether potential alternatives achieve comparable results, *e.g.*, whether a policy accurately predicts the risk of loan default while producing less discriminatory outcomes.

11. BLDS's proprietary analyses thus track the elements of traditional disparate-impact doctrine to assist clients in complying with it. The market for it is intrinsically bound up with the traditional doctrine. That market will dry up or disappear entirely under the framework advanced in HUD's rule. For example, if defendants can prevail in litigation regardless of whether they can accomplish their legitimate objectives through less discriminatory alternatives, lenders and other market participants will be unlikely to retain BLDS to assess whether these models cause adverse disparate impacts or whether less discriminatory

alternatives exist.

12. HUD's rule, if permitted to remain in effect, thus will directly reduce or destroy altogether a significant and consistent existing source of revenue for BLDS. It also will directly reduce the value of investments that BLDS has made into the analysis of machine-learning models that otherwise are likely to produce significantly more revenue in the future.

13. In recent years, machine-learning models have introduced a new form of complexity to the disparate-impact analysis. These models assess the effects of a large number of potential variables—not simply individually, but in various combinations—and determine which combinations best predict outcomes such as likelihood of loan defaults. Entities are increasingly using artificial intelligence models to make decisions regarding creditworthiness, marketing, and other key issues related to housing.

14. Machine-learning models can be very predictive and offer advantages over traditional credit models. Indeed, machine-learning models often can be good alternatives to traditional models. They can have similar or superior predictive quality but, because they rely on different variables and are susceptible to a range of alterations, they can have very different levels of discriminatory effect. These models are also often able to rely on a larger variety of data, beyond those used in traditional models and can result in making more credit available to racial minorities—who, for example, regularly pay their bills but have credit histories that are too "thin" to be given high credit scores now.

15. Machine-learning models, however, can have a black-box quality that makes it difficult to determine why they are causing a disparate impact. They may rely on many variables, some of which would not be obvious choices for a human but that nonetheless correlate with risk. Adding further complexity, those variables often correlate to risk only in

complex combinations that, again, are not obvious for a human.

16. To demonstrate the sort of variables that can feature in such models, one recent study showed that whether a person uses an iPhone vs. an Android smartphone is as predictive of credit risk as a large difference in credit score.[1] Almost certainly, it is not the smartphone that is the real cause of greater or lesser credit risk; rather, the model is finding information that happens to correlate with things that *do* make someone a better or worse credit risk but are not publicly available. But the type of smartphone used *also* correlates to some extent with race, and so reliance on this variable will introduce a discriminatory effect. And frequently a model will find and rely upon complex combinations of such non-intuitive variables, making it difficult to know how dependent the model's results are on any of them.

17. All of this is to say that a machine-learning model can easily find many patterns that correlate in some respect to risk but that also cause disparate impacts based on race or other protected class. There thus is an obvious need for sophisticated analysis of these models to ensure that they comply with the traditional disparate-impact doctrine and do not have unnecessary discriminatory effects. That is, can alternative predictive combinations of variables be found that are predictive but that have less of an adverse correlation with race or other protected classes. But HUD's proposed rule negates the need for that analysis by effectively immunizing policies and practices—including the use of models—from traditional disparate-impact scrutiny.

18. The steps in the traditional disparate impact analysis—assessing whether a policy or practice causes a disparate adverse effect, assessing whether that policy furthers a

[1] Tobias Berg et al., *On the Rise of the FinTechs—Credit Scoring Using Digital Footprints* 3-4 (Fed. Deposit Ins. Corp. Ctr. for Fin. Research, Working Paper No. 2018-04), https://perma.cc/RAZ6-VPXX.

legitimate business reason, and assessing whether less discriminatory alternatives exist—have been essential under existing disparate impact law and are a core component of BLDS's work on behalf of clients. HUD's rule so drastically weakens the disparate impact doctrine that it makes these analyses irrelevant with respect to entities' policies and practices, including their use of models.

19. Assessing whether models cause disparate impact, advance legitimate business interests, and are the least discriminatory options are not insurmountable obstacles. These are technical problems that have technical solutions.

20. In the last few years, BLDS has devoted considerable resources to developing proprietary methods for analyzing traditional and machine-learning models to determine whether these models unnecessarily cause discriminatory outcomes and then altering these models to reduce that discriminatory effect while preserving their predictive quality in an efficient and cost-effective way.

21. These proprietary methods for reducing the discriminatory effect of models are of considerable value under the traditional disparate-impact rules. Indeed, companies that perform similar services, albeit not as efficient or cost effective, are being valued at millions of dollars by investors. For example, our methodology is implemented in an easy to use and cost-effective software package that would be of great use to regulators and to every creditor in America using automated models. Right now, we anticipate this being our core growth area, which is why we have invested so much time and money into developing methods that make us industry leaders.

22. If HUD's rule is permitted to go into effect, the regulatory environment would radically change, and the value of these proprietary methods would be much reduced. Lenders

and other providers of housing-related services would face much less risk for using algorithms with an unnecessary disparate impact rather than adopting less discriminatory alternatives.

23. HUD's proposed rule would have provided a defense to disparate impact claims that provided immunity for entities that showed that a model did not rely on factors that were substitutes or close proxies for protected classes under the Fair Housing Act and that the model was predictive of credit risk or other similar valid objectives. As commenters noted, that proposed defense was misguided and ill-conceived for a number of reasons, including that the proposed defense would have allowed that entity to escape liability even if a model caused a disparate impact and less discriminatory alternatives were available. HUD did not finalize that proposed defense.

24. However, HUD did finalize a different new defense that effectively immunizes a defendant that demonstrates that a policy or practice "is intended to predict an occurrence of an outcome, the prediction represents a valid interest, and the outcome predicted by the policy or practice does not or would not have a disparate impact on protected classes compared to similarly situated individuals not part of the protected class."

25. Because HUD worded this defense so ambiguously, and because it did not offer the public the opportunity to comment on it, it is unclear what this defense means or how this defense should be applied. BLDS will be required to spend time and resources assessing this new defense and if possible designing or modifying analyses and methodologies to account for this new defense.

26. HUD explains the logic behind it in the Supplementary Information to the Final Rule. HUD asserts that if:

> a plaintiff alleges that a lender rejects members of a protected class at higher rates than non-members, then the logical conclusion of such claim would be that

> members of the protected class who were approved, having been required to meet an unnecessarily restrictive standard, would default at a lower rate than individuals outside the protected class. Therefore, if the defendant shows that default risk assessment leads to less loans being made to members of a protected class, but similar members of the protected class who did receive loans actually default more or just as often as similarly situated individuals outside the protected class, then the defendant could show that the predictive model was not overly restrictive.

27. This argument is incorrect and would allow clearly discriminatory decisions to be acceptable. By focusing only the "outcome" predicted by a policy or practice (e.g., default), the defense only considers the characteristics of individuals that *benefited* from a policy or practice (e.g., those that were "approved" for loans and therefore would have a possibility of experiencing a default outcome). It ignores, without any basis, the characteristics of individuals that applied but were *excluded* by a policy or practice. Without considering the characteristics of all applicants, it is not possible to determine that the selection process was more lenient or "overly restrictive" for a protected class because it ignores the incidence of characteristics among rejected applicants. It also ignores features that often vary *within* the accepted pool and that contribute to default rates: for example, increased costs and fees, higher interest rates, or less permissive waiver policies. In other words, as a statistical matter, it is simply not true that equal or higher default rates among protected class members indicate whether a decision process was "overly restrictive" or discriminatory.

28. I am unaware of any court or regulator approving of such a defense. I am unaware of any entity employing this type of analysis in its fair lending compliance programs.

29. I am aware of at least one employment discrimination suit where a court rejected arguments akin to the HUD defense. *See Com. of Pa. v. O'Neill*, 348 F. Supp. 1084, 1095-96 (E.D. Pa. 1972), *order vacated in part*, No. 72-1614, 1972 WL 2595 (3d Cir. Sept. 14, 1972), *on reh'g*, 473 F.2d 1029 (3d Cir. 1973), *and af''d in part, vacated in part*, 473 F.2d 1029

(3d Cir. 1973). Defendant in that suit attempted to defend against disparate impact claims by focusing only on characteristics of employees that benefitted from a policy or practice (i.e., "accepted" employees), and argued that characteristics of employees only within that "accepted" population indicated the defendant must have been more lenient in accepting minority applicants. The court rejected that defense as unsound, relying in part on my own expert analysis.

30. The defense has other problems as well. With respect to algorithms (e.g., an automated credit score model), if the term "similarly situated" as used in the new defense is interpreted as scoring the same in the algorithm, the defense would be an assessment of model bias and validity—concepts that are separate and independent from established disparate impact analysis—and it would eliminate the requirement to adopt a model with less disparate impact and comparable performance.

31. By promulgating this new disparate impact defense, the Final Rule will create confusion among institutions seeking to comply with fair housing and lending laws, plaintiffs, and courts. It will also shield discriminatory practices, and for algorithmic models, it will eliminate the need to adopt high-performing alternative models that have less disparate impact.

32. With respect to algorithms, some entities will view this new defense as an avenue for establishing immunity from disparate impact claims, thereby eliminating their litigation and regulatory risks. That view will diminish demand for an important part of our practice and significantly reduce the value of some of our intellectual property.

33. I am very familiar with the way that regulatory and legal change predictably affects the demand for our services and the value of our expertise and proprietary methods. We regularly have discussions with clients about the regulatory environment and litigation risk, as

they decide whether they need specific services. Clients and prospective clients are very clear that demand for our services decreases if the clients perceive a decrease in litigation or regulatory risk. For example, I once had a thriving practice advising companies regarding the unnecessary disparate impact caused by their corporate practices notwithstanding that they had delegated decision-making authority to various lower level supervisors. After the Supreme Court's decision in *Wal-Mart v. Dukes*, 564 U.S. 338 (2011), the demand for that service dropped off dramatically. That is because many companies no longer believed they faced significant litigation risk for maintaining corporate practices with measurable disparate impact so long as final decisions are made at a lower level. It is foreseeable that HUD's rule will have the same effect, and we expect that it will lead to a significant downturn in our current work as well the loss of a major additional revenue stream in the future. Indeed, it appears that HUD's rule is intended to have that effect, *i.e.*, to make it unnecessary for companies to conduct the sort of analysis of their algorithms, policies, and practices that BLDS offers.

34. The bottom line is that, if this rule goes into effect, it will directly diminish demand for an important part of our practice and significantly reduce the value of some of our intellectual property.

I declare under penalty of perjury that the foregoing is true and correct.

EXECUTED WITHIN THE UNITED STATES ON: October 16, 2020.

Bernard R. Siskin

BY: ______________________________
Dr. Bernard Siskin